On the Capacity Advantage of a Single Bit

Michael Langberg Michelle Effros

arXiv:1607.07024v1 [cs.IT] 24 Jul 2016

***Abstract*—In this work we study the capacity advantage achieved by adding a *single bit* of communication – not a link of capacity 1 but a single bit over all time – to a memoryless network. Specifically, we present a memoryless network in which adding a single bit of communication strictly increases the capacity region.**

I. Introduction

The *edge removal problem*, defined and studied in [1], [2], aims to quantify the loss in capacity that results from the removal of a single edge (i.e., a point-to-point channel) from a given (possible noisy) communication network. At first glance, it may seem that the removal of an edge with edge capacity $\delta > 0$ should change the network capacity by an amount that tends to zero as δ tends to zero. The following examples demonstrate that this is not always the case.

In [3], [4], Guruswami and Langberg present simple examples of point-to-point channels with memory for which the capacity of a network $\mathcal{N}$ containing the channel and a *side-information* edge of capacity δ from the transmitter to the receiver is described by a function $C_{\mathcal{N}}(\delta)$ that is not continuous at $\delta = 0$.

While the discontinuity of $C_{\mathcal{N}}(\delta)$ in [3], [4] can be attributed to the memory in the point-to-point channel, [5] shows a similar phenomenon for a memoryless network. Specifically, [5] studies communication over a network $\mathcal{N}$ containing a memoryless Multiple Access Channels (MAC) and a δ-capacity *cooperation* edge carrying information from a "cooperation facilitator" to the two transmitters. Once again, the capacity $C_{\mathcal{N}}(\delta)$ exhibits a discontinuity at $\delta = 0$. More specifically, [5] shows that for any MAC whose average- and maximal-error capacities differ, adding a cooperation facilitator with output capacity δ results in a network $\mathcal{N}$ whose maximal-error capacity $C_{\mathcal{N}}(\delta)$ exhibits the described discontinuity. Dueck's 2-way *contraction channel* [6] is an example of a MAC with the described property.

In this work, we take the edge removal problem to an extreme, seeking to understand the effect of removing not a δ-capacity edge for arbitrarily small $\delta > 0$ but rather an edge that can only carry 1 bit of communication (over all time). We wish to understand whether there exist networks for which the removal of a single bit of communication can strictly change the capacity region. Our study focuses on noisy networks. We discuss noiseless networks, i.e., network coding, at the end of this section.

M. Langberg is with the Department of Electrical Engineering at The State University of New York at Buffalo. Email : mikel@buffalo.edu

M. Effros is with the Department of Electrical Engineering at the California Institute of Technology. Email : effros@caltech.edu

This material is based upon work supported by NSF grants CCF-1018741 and CCF-1016671.

For noisy channels with memory, the answer is immediate. It is not difficult to construct a two-state point-to-point channel with memory for which a single bit of feedback changes the network capacity. For example, consider a binary symmetric channel whose error probability θ is chosen at random and then fixed for all time. If θ equals θ_i, $i \in \{1, 2\}$, with positive probability p_i and $0 \leq \theta_1 < \theta_2 \leq 1/2$ then a single bit of feedback from the receiver to the transmitter asymptotically suffices to increase the capacity.

For memoryless channels, the question is far more subtle. Its solution is the subject of this work. In Section II, we formalize the notion of a single bit of communication by defining the "1-bit channel" and the corresponding notion of capacity to accommodate networks with 1-bit channels. In Section III, we demonstrate the existence of a network for which the removal of a single 1-bit channel changes the capacity region significantly. We employ the contraction channel of Dueck [6] in our construction.

We note that the edge removal problem has seen significant studies in the noiseless setting of network coding as well. In contrast to the noisy communication setting, in network coding instances, the question whether there exists networks for which the removal of edges of negligible capacity (in the block length n) causes a strictly positive loss in rate is an open problem. For several network coding instances studied in [1], [2] it is shown that the removal of an edge of capacity δ may decrease the rate of communication between each source-receiver pair by at most δ. These instances include networks with collocated sources, networks in which we are restricted to perform linear encoding, networks in which the edges removed are connected to receivers with no out going edges, and additional families of network coding instances. However, whether there exist network coding instances with a δ capacity edge for which the capacity region $C(\delta)$ is not continuous at $\delta = 0$ remains an intriguing open problem connected to a spectrum of (at times seemingly unrelated) questions in the context of network communication, e.g., [7], [8], [9], [10], [11], [12], [13]. One can study the effect of removing 1-bit channels in the noiseless setting of network coding. It is tempting to believe that the removal of a 1-bit channel cannot effect the capacity region of network coding instances, however, the problem is left open in this work.

II. Model

We wish to consider memoryless networks enhanced by the addition of one or more 1-bit channels. We begin by defining memoryless and 1-bit channels and then define codes and capacities appropriate for networks that combine them.

The following definitions are useful in the description that follows. For any positive constant c, let $[c] = \{1, \ldots, \lfloor c \rfloor\}$. For any sets $A \subseteq B$ and any values x_i or sets $\mathcal{X}_i$ indexed by $i \in B$, let $x_A = (x_i : i \in A)$ and $\mathcal{X}_A = \prod_{i \in A} \mathcal{X}_i$.

A. Networks

An m-node *memoryless network* is described by a triple

$$\left(\mathcal{X}_{[m]}, p(y_{[m]}|x_{[m]}), \mathcal{Y}_{[m]}\right),$$

where $x_i \in \mathcal{X}_i$ and $y_i \in \mathcal{Y}_i$ represent the network's inputs from and outputs to node i. Thus, at each time t node i transmits a network input $X_{i,t}$ and receives a network output $Y_{i,t}$ governed by the statistical relationship $p(y_{[m]}|x_{[m]})$ relating all network inputs to all network outputs at time t.

A *1-bit channel* is a point-to-point channel from some node i to another node j in the network. Unlike the memoryless network above, which can carry the same amount of information in every time step, the 1-bit channel can carry only one bit of information in total. That information is transmitted at some fixed time τ chosen in the code design. Given τ, the 1-bit channel effectively acts like a collection of independent, single-time-step memoryless channels, where the channel at time t is

$$\begin{array}{ll} (\{0,1\}, 1(y_{j,t} = x_{i,t}), \{0,1\}) & \text{if } t = \tau \\ (\{0\}, 1(y_{j,t} = x_{i,t}), \{0\}) & \text{if } t \neq \tau. \end{array}$$

Here we use an alphabet $\{0\}$ of size 1 to represent the fact that the 1-bit channel is inactive when $t \neq \tau$.

B. Codes

Consider a network $\mathcal{N}$ that combines a memoryless channel with one or more 1-bit channels. To capture the 1-bit channels, we hence forward allow alphabets $\mathcal{X}_{i,t}$ and $\mathcal{Y}_{i,t}$ to vary with t since node i may transmit information to or receive information from one or more 1-bit channels.

A *blocklength-n code* $(\tau, X, \hat{M})$ operates the network $\mathcal{N}$ over n time steps with the goal of transmitting a message $M_i \in \mathcal{M}_i$ from each node $i \in [m]$ to all receivers j in demand set $D_i \subseteq [m] \setminus \{i\}$. By assumption, $\mathcal{M}_i = [2^{nR_i}]$, R_i equals zero if and only if $D_i = \emptyset$, and $M_{[m]}$ is distributed uniformly on $\mathcal{M}_{[m]}$. Code $(\tau, X, \hat{M})$ is defined by its activation vector

$$\tau = (\tau_e)_{e \in \alpha},$$

encoding functions

$$X = (X_{i,t})_{i \in [m], t \in [n]},$$

and decoding functions

$$\hat{M} = (\hat{M}_{j,i})_{j \in [m], i \in D_j}.$$

Here, α represents the set of 1-bit channels in $\mathcal{N}$. For each 1-bit channel e, constant $\tau_e \in [n]$ describes the activation time of e. For each i and t, encoder

$$X_{i,t} : \mathcal{Y}_{i,[t-1]} \times \mathcal{M}_i \to \mathcal{X}_{i,t}$$

maps node i's previously received network outputs and outgoing message to the time-t, node-i network input $X_{i,t}$. For each message M_j and node i that demands M_j, the decoder

$$\hat{M}_{j,i} : \mathcal{Y}_{i,[n]} \times \mathcal{M}_i \to \mathcal{M}_j$$

maps node i's received network outputs and outgoing message to a reconstruction of message M_j.

C. Capacities

For the purpose of defining capacity, it is useful to characterize a code $(\tau, X, \hat{M})$ by its blocklength n, rate $R_{[m]}$, and error probability. The literature contains several definitions of error probability, most notably the *maximal error probability*

$$d_{\max}(\tau, X, \hat{M}) = \max_{j \in [m], i \in D(j)} \Pr(\hat{M}_{j,i} \neq M_j)$$

and the *average error probability* $d_{\text{avg}}(\tau, X, \hat{M})$:

$$\frac{1}{|\mathcal{M}_{[m]}|} \sum_{j \in [m], w_j \in \mathcal{M}_j} \Pr(\exists i \in D_j : \hat{M}_{j,i} \neq w_j \mid M_j = w_j).$$

A blocklength-n, rate-$R_{[m]}$ code $(\tau, X, \hat{M})$ is called a max-$(\varepsilon, R_{[m]}, n)$ solution if $d_{\max}(\tau, X, \hat{M}) \leq \varepsilon$ and an avg-$(\bar{\varepsilon}, R_{[m]}, n)$ solution if $d_{\text{avg}}(\tau, X, \hat{M}) \leq \bar{\varepsilon}$.

Definition 2.1 (Maximal-Error Capacity): The maximal-error capacity region $\mathcal{C}_{\max}(\mathcal{N})$ of network $\mathcal{N}$ is the closure of all rate vectors $R_{[m]}$ such that for any $\varepsilon > 0$, there exists a max-$(\varepsilon, R_{[m]}, n)$ solution for all n sufficiently large.

Definition 2.2 (Average-Error Capacity): The average-error capacity region $\mathcal{C}_{\text{avg}}(\mathcal{N})$ of network $\mathcal{N}$ is the closure of all rate vectors $R_{[m]}$ such that for any $\bar{\varepsilon} > 0$, there exists an avg-$(\bar{\varepsilon}, R_{[m]}, n)$ solution for all n sufficiently large.

D. Remarks

Some remarks are in order. Notice that for $R^{[m]}$ to be included in (the interior of) a capacity region $\mathcal{C}_{\max}(\mathcal{N})$ (similarly for $\mathcal{C}_{\text{avg}}(\mathcal{N})$), we require the existence of max-$(\varepsilon, R^{[n]}, n)$ codes $(\tau, X, \hat{M})$ for *all* sufficiently large blocklengths. Without such a requirement, the addition of 1-bit channels can have a significant effect on the achievable rate for small blocklengths n. Take for example an empty network $\mathcal{N}$ enhanced with a single 1-bit channel e to obtain $\mathcal{N}^+$. For $n = 1$, there is a clear difference between the rates achievable on $\mathcal{N}$ and $\mathcal{N}^+$. We therefore require the blocklength n to grow without bound to capture the idea that the channel e can carry only one bit *over all time* rather than that e can carry one bit in a small time window $[n]$. Notice further that our definitions require each 1-bit channel $e \in \alpha$ to be active at a constant time τ_e independent of the messages $M_{[m]}$. This is important since it prevents codes that might use timing to convey information about the messages.

III. Main result

The main question we ask in this work is whether there exists a network $\mathcal{N}$ and a 1-bit channel e such that adding e to $\mathcal{N}$ yields a new network $\mathcal{N}^*$ with a strictly larger capacity region. Namely, with

$$\mathcal{C}(\mathcal{N}^*) \neq \mathcal{C}(\mathcal{N}).$$

We answer the question for maximal error below. We leave the question in the context of average error open in this work.[1]

To prove our result, we start by constructing a network $\mathcal{N}^+$ by combining Dueck's memoryless MAC [6] with a pair of memoryless point-to-point channels and a collection α of 1-bit channels. The construction, described below, is depicted in Figure 1.

Dueck's MAC has a pair of transmitters, nodes 1 and 2, and a single receiver, node 3. The input and output alphabets are

$$\begin{array}{lllllll} \mathcal{X}_1 & = & \{a,b,A,B\} & \mathcal{Y}_1 & = & \{0\} \\ \mathcal{X}_2 & = & \{0,1\} & \mathcal{Y}_2 & = & \{0\} \\ \mathcal{X}_3 & = & \{0\} & \mathcal{Y}_3 & = & \{a,b,c,A,B,C\} \times \{0,1\} \end{array}$$

Since nodes 1 and 2 receive no outputs from the channel and node 3 has no input to the channel (here denoted by alphabets of size 1) we simplify the notation from Section II to describe the channel as $\mathcal{N}_D = (\mathcal{X}_1 \times \mathcal{X}_2, p(y_3|x_1,x_2), \mathcal{Y}_3)$. Where $p(y_3|x_1,x_2) = 1(y_3 = W(x_1,x_2))$ captures the channel's deterministic behavior with output

$$W(x_1,x_2) = \begin{cases} (c,0) & \text{if } (x_1,x_2) \in \{(a,0),(b,0)\} \\ (C,1) & \text{if } (x_1,x_2) \in \{(A,1),(B,1)\} \\ (x_1,x_2) & \text{otherwise.} \end{cases}$$

We next add to Dueck's MAC $\mathcal{N}_D$ a fourth node, here called the "cooperation facilitator" (CF), and a pair of memoryless, point-to-point channels. We denote the resulting network by $\mathcal{N}_0$. The first memoryless channel, from node 1 to node 4, is a noiseless channel of capacity 2. The second, from node 2 to node 4, is a noiseless channel of capacity 1. Given a pair of messages originating at nodes 1 and 2 and a single receiver at node 3 ($D_1 = D_2 = \{3\}$ and $D_3 = D_4 = \emptyset$), the capacity region of this modified channel and Dueck's MAC are identical ($\mathcal{C}_{\max}(\mathcal{N}_0) = \mathcal{C}_{\max}(\mathcal{N}_D)$).

Finally, we build network $\mathcal{N}^+$ from memoryless network $\mathcal{N}_0$ by adding k 1-bit channels, $\{f_1,\ldots,f_k\}$, from the CF (node 4) to receiver node 3, and by adding a single 1-bit channel, e, from the CF to node 1.

Our main result shows that for some finite k,

$$\mathcal{C}_{\max}(\mathcal{N}^+) \neq \mathcal{C}_{\max}(\mathcal{N}_0).$$

That is, removing the 1-bit edges leaving node CF changes the capacity region of the network.

Theorem 3.1: For some finite constant k,

$$(1.24, 1) \in \mathcal{C}_{\max}(\mathcal{N}^+)$$

[1]The results of [5] that exhibit capacity $C_{\mathcal{N}}(\delta)$ with a discontinuity at $\delta = 0$ hold for maximal error only. The average error case remains open.

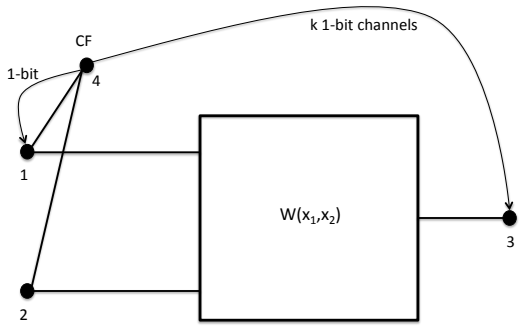

Fig. 1: Network $\mathcal{N}^+$. Transmitter nodes 1 and 2 wish to send a pair of independent messages to the receiver node 3. The cooperation facilitator, node 4, is labeled by CF. Node 1 receives a 1-bit input from the CF. Receiver node 3 receives an additional k-bit input from the CF.

while,

$$(1.19, 0.97) \notin \mathcal{C}_{\max}(\mathcal{N}_0)$$

As a corollary of Theorem 3.1 we have:

Corollary 3.1: There exists a network $\mathcal{N}$ and a single 1-bit channel e such that the network $\mathcal{N}^*$ obtained by adding e to $\mathcal{N}$ satisfies:

$$\mathcal{C}_{\max}(\mathcal{N}^*) \neq \mathcal{C}_{\max}(\mathcal{N}).$$

Proof: (of Corollary 3.1) Consider adding the $(k+1)$ 1-bit channels of $\mathcal{N}^+$ one by one to the network $\mathcal{N}_0$. Specifically, let $\mathcal{N}_1$ be the network $\mathcal{N}_0$ enhanced by one of the 1-bit channels of $\mathcal{N}^+$ (which channel precisely is of no significance to the current proof; it may be chosen arbitrarily). Similarly, for $i = 2,\ldots,k+1$, let $\mathcal{N}_i$ be the network $\mathcal{N}_{i-1}$ enhanced by one of the 1-bit channels of $\mathcal{N}^+$ that do not yet appear in $\mathcal{N}_{i-1}$. It holds that

$$\mathcal{C}_{\max}(\mathcal{N}_0) \subseteq \mathcal{C}_{\max}(\mathcal{N}_1) \subseteq \mathcal{C}_{\max}(\mathcal{N}_2) \cdots \subseteq \mathcal{C}_{\max}(\mathcal{N}_{k+1}).$$

Since $\mathcal{C}_{\max}(\mathcal{N}_0) \neq \mathcal{C}_{\max}(\mathcal{N}^+)$ by Theorem 3.1, there must be an i for which $\mathcal{C}_{\max}(\mathcal{N}_i) \neq \mathcal{C}_{\max}(\mathcal{N}_{i+1})$. The corollary then follows by setting $\mathcal{N}$ equal to $\mathcal{N}_i$ and e to be the 1-bit channel added to $\mathcal{N}_i$ to obtain $\mathcal{N}_{i+1}$. ■

A. Proof of Theorem 3.1

First, we note that it is proven in [6] that for $p \in [0,1/2]$

$$(H(1/3) + 2/3 - p, H(p)) \notin \mathcal{C}_{\max}(\mathcal{N}_D) = \mathcal{C}_{\max}(\mathcal{N}_0).$$

Setting $p = 0.4$, which gives $H(p) = 0.97$, we note that $(1.19, 0.97) \notin \mathcal{C}_{\max}(\mathcal{N}_0)$.

Let $\delta > 0$ be a sufficiently small constant and let $k = \lceil \log(8/\delta) \rceil + 1$ (so that $\mathcal{N}^+$ adds $\lceil \log(8/\delta) \rceil + 2$ 1-bit channels to $\mathcal{N}_0$). We next prove that

$$(1.25 - 0.5\delta, 1) \in \mathcal{C}_{\max}(\mathcal{N}^+)$$

by demonstrating the existence, for any sufficiently large n, of a zero-error $(2n+1)$-block length encoding scheme for network $\mathcal{N}^+$ of rate asymptotically close to $(1.25-0.5\delta, 1)$. We employ a coding scheme that describes the message sets $\mathcal{M}_1$ and $\mathcal{M}_2$ in two parts, here denoted by $\mathcal{M}_1 = \mathcal{M}_1^1 \times \mathcal{M}_1^2 = [2^n] \times [2^{\frac{3n}{2}-\delta n}]$, and $\mathcal{M}_2 = \mathcal{M}_2^1 \times \mathcal{M}_2^2 = [2^n] \times [2^n]$. The first part of each message is described in the first n time steps. All 1-bit channels are activated at time $n+1$. The second part of each message is then described at time steps $n+2$ through $2n+1$.

A rough description of our scheme follows. In the first phase messages $M_1^1 \in \mathcal{M}_1^1$ and $M_2^1 \in \mathcal{M}_2^1$ are communicated through the network at rate $(1,1)$. Transmitter node 1 and receiver node 3 cannot use in their encoding/decoding functions information from their incoming 1-bit channels since the 1-bit channels have not been activated yet. During this first phase, the codewords corresponding to messages M_1^2 and M_2^2 are forwarded on the edges $(1, CF)$ and $(2, CF)$ to the cooperation facilitator CF. At time step $n+1$ the CF, knowing M_1^2 and M_2^2, computes and sends the outputs to the 1-bit channels. Finally, during time steps $n+2$ to $2n+1$, messages M_1^2 and M_2^2 are communicated through the network. Our scheme in the second phase is described in detail below. The strategy strongly relies on the information received at transmitter node 1 and receiver node 3 from the 1-bit channels leaving the CF. The communication rate in the second phase is $(3/2-\delta, 1)$. All in all, the total rate over both phases tends to $(1.25-0.5\delta, 1)$ (as n tends to infinity).

1) Phase 1 (time steps 1 to n): In phase 1, the codebook for $\mathcal{M}_1^1$ equals the set of all vectors in $\{a, A\}^n$. The codebook for $\mathcal{M}_2^1$ equals the set of all the vectors in $\{0,1\}^n$. It is not hard to verify that for any $x_1 \in \{a, A\}$ and $x_2 \in \{0,1\}$ one can decode (with zero error) the value of x_1 and x_2 from $W(x_1, x_2)$. This implies a rate of $(1,1)$ in the first n time steps. In addition to communicating information regarding M_1^1 and M_2^1, in this first phase, the n time steps are used to forward the codewords for messages M_1^2 and M_2^2 (to be specified below) from transmitter nodes 1 and 2 to the CF.

2) 1-bit channel activation phase (time step $n+1$): The 1-bit channels in $\mathcal{N}^+$ are all activated at time step $n+1$. Up to time $n+1$, the CF receives the codewords for messages M_1^2 and M_2^2 from transmitter nodes 1 and 2, and may compute the bits to be transmitted on its outgoing 1-bit channels. The codewords for messages M_1^2 and M_2^2 and the functions to be computed for each outgoing 1-bit channel are specified in the description of phase 2.

3) Phase 2 (time steps $n+2$ to $2n$): In phase 2, the codebook for $\mathcal{M}_2^2$ (as before) contains all vectors in $\{0,1\}^n$. In this phase, the codebook for transmitter node 1 contains $2|\mathcal{M}_1^2|$ codewords. We here multiply the size of $\mathcal{M}_1^2$ by 2 as transmitter node 1 may use in its encoding process not only the input message M_1^2 but also the input bit b received at time step $n+1$ from the CF over the 1-bit channel. The first $|\mathcal{M}_1^2|$ codewords for transmitter node 1 are a set of distinct vectors from $\mathcal{X}_1^n$, here denoted by $\{\mathbf{x}_1(w)\}$ where $w = 1, \ldots, |\mathcal{M}_1^2|$. The set $\{\mathbf{x}_1(w)\}_w$ is chosen uniformly at random from the collection of all subsets of $\mathcal{X}_1^n$ of size $|\mathcal{M}_1^2|$, i.e., from $\{A \subset \mathcal{X}_1^n \mid |A| = |\mathcal{M}_1^2|\}$. The second $|\mathcal{M}_1^2|$ codewords for transmitter node 1 are the codewords $\bar{\mathbf{x}}_1(w)$ for $w = 1, \ldots, |\mathcal{M}_1^2|$ where $\bar{\mathbf{x}}_1(w)$ is a *toggled* version of $\mathbf{x}_1(w)$ in which each lowercase entry of $\mathbf{x}_1(w)$ is converted to an uppercase value (of the same letter) and vica versa. Given message M_1^2, node 1 transmits $\mathbf{x}_1(M_1^2)$ if it receives $b=0$ on its 1-bit channel and transmits $\bar{\mathbf{x}}_1(M_1^2)$ if it receives $b=1$ on its 1-bit channel.

As specified previously, during phase 1 of our communication, the codewords corresponding to M_1^2 and M_2^2 are transmitted to the CF. Specifically, node 1 transmits $\mathbf{x}_1(M_1^2)$ to the CF and node 2 transmits $\mathbf{x}_2(M_2^2)$ to the CF. Before defining the encoding function of the CF and the decoding function of receiver node 3 we set some notation.

Let $\mathcal{Y}_3 = \mathcal{Y}_{3,1} \times \mathcal{Y}_{3,2} = \{A, B, C, a, b, c\} \times \{0,1\}$ be the output alphabet at receiver node 3. Let $\mathbf{y}_{3,1} = (y_{3,1;1}, \ldots, y_{3,1;n}) \in \mathcal{Y}_{3,1}^n$ and $\mathbf{y}_{3,2} = (y_{3,2;1}, \ldots, y_{3,2;n}) \in \mathcal{Y}_{3,2}^n$ be the n symbols received in phase 2 at node 3. Specifically, given transmitted codewords $\mathbf{x}_1$ and $\mathbf{x}_2$, we have $(\mathbf{y}_{3,1}, \mathbf{y}_{3,2}) = W^n(\mathbf{x}_1, \mathbf{x}_2)$, where W^n is the blocklength-n extension of W. We use notation $\mathbf{y}_{3,1} = W_1^n(\mathbf{x}_1, \mathbf{x}_2)$ and $\mathbf{y}_{3,2} = W_2^n(\mathbf{x}_1, \mathbf{x}_2)$, where $W(x_1, x_2) = (W_1(x_1, x_2), W_2(x_1, x_2))$.

Note that $W_2(\mathbf{x}_1, \mathbf{x}_2) = \mathbf{x}_2$ for any $\mathbf{x}_1$ and $\mathbf{x}_2$, thus $W^n(\mathbf{x}_1, \mathbf{x}_2) = (\mathbf{y}_{3,1}, \mathbf{x}_2)$. Let $C(\mathbf{y}_{3,1}) = |\{i \mid \mathbf{y}_{3,1;i} \in \{c, C\}\}|$ be the number of symbols in $\mathbf{y}_{3,1}$ that equal c or C. If $\mathbf{X}_1(\mathbf{y}_{3,1}, \mathbf{x}_2)$ is the set of elements $\mathbf{x}_1 \in \mathcal{X}_1^n$ for which $W^n(\mathbf{x}_1, \mathbf{x}_2) = (\mathbf{y}_{3,1}, \mathbf{x}_2)$ then $|\mathbf{X}_1(\mathbf{y}_{3,1}, \mathbf{x}_2)| = 2^{C(\mathbf{y}_{3,1})}$. As one can decode $\mathbf{x}_2$ given $\mathbf{y}_{3,1}$ (this follows directly from the definition of W), we henceforth use the notation $\mathbf{X}_1(\mathbf{y}_{3,1})$ instead of $\mathbf{X}_1(\mathbf{y}_{3,1}, \mathbf{x}_2)$ respectively.

We say that $\mathbf{y}_{3,1}$ is *good* if $|\mathbf{X}_1(\mathbf{y}_{3,1})| \leq 2^{n/2}$. In our encoding scheme for phase 2, we would like to guarantee that $\mathbf{y}_{3,1}$ is always good. This is accomplished using the information b from the 1-bit channel sent between the CF and node 1. Namely, let $\mathbf{x}_1$ and $\mathbf{x}_2$ be any pair of codewords transmitted to the CF during phase 1 of our communication. Given $\mathbf{x}_1$ and $\mathbf{x}_2$, the CF may compute $W^n(\mathbf{x}_1, \mathbf{x}_2) = (\mathbf{y}_{3,1}, \mathbf{x}_2)$. If $\mathbf{y}_{3,1}$ is good, the CF sends the bit $b=0$ to transmitter node 1 at time step $n+1$; otherwise the CF sends the bit $b=1$. Now node 1, knowing bit b, sends codeword $\mathbf{x}_1$ over the n time steps of phase 2 if $b=0$ and $\bar{\mathbf{x}}_1$ if $b=1$.

Let $\mathbf{x}_1(w)$ and $\bar{\mathbf{x}}_1(w)$ be the codewords corresponding to message $w = M_1^2$ of transmitter node 1 and let $\mathbf{x}_2(v)$ be the codeword corresponding to message $v = M_2^2$ of transmitter node 2. If the bit sent from the CF to transmitter node 1 in time step $n+1$ is 0, then the received symbols at node 3 are $(\mathbf{y}_{3,1}, \mathbf{x}_2(v)) = W^n(\mathbf{x}_1(w), \mathbf{x}_2(v))$, and if the bit sent from the CF to transmitter node 1 in time step $n+1$ is 1, then the received symbols at node 3 are $(\mathbf{y}_{3,1}, \mathbf{x}_2(v)) = W^n(\bar{\mathbf{x}}_1(w), \mathbf{x}_2(v))$.

Claim 3.1: For any $\mathbf{x}_1(w)$ and $\mathbf{x}_2(v)$, if $W_1^n(\mathbf{x}_1(w), \mathbf{x}_2(v))$ is not good then $W_1^n(\bar{\mathbf{x}}_1(w), \mathbf{x}_2(v))$ is good.

Proof: Follows from the definition of W^n. Consider an entry x_1 in $\mathbf{x}_1(w)$ and its corresponding entry $\bar{x}_1$ in

$\bar{\mathbf{x}}_1(w)$. Let $x_2 \in \{0,1\}$. It holds (by a simple exhaustive case analysis) that exactly one of the values $W_1(x_1,x_2)$ and $W_1(\bar{x}_1,x_2)$ is in the set $\{c,C\}$. Thus, the number of entries in $W_1^n(\mathbf{x}_1(w),\mathbf{x}_2(v))$ which are in the set $\{c,C\}$ plus the number of entries in $W_1^n(\bar{\mathbf{x}}_1(w),\mathbf{x}_2(v))$ which are in the set $\{c,C\}$ is exactly n. If the former is more than $n/2$ (implied by the fact that $W_1^n(\mathbf{x}_1(w),\mathbf{x}_2(v))$ is not good) then the latter is less than $n/2$ (implying that $W_1^n(\bar{\mathbf{x}}_1(w),\mathbf{x}_2(v))$ is good). ■

We now show that with high probability over our code design, for any good received word $(\mathbf{y}_{3,1},\mathbf{y}_{3,2})$, there are at most $\lceil 8/\delta \rceil$ codeword pairs $(\mathbf{x}_1,\mathbf{x}_2)$ that satisfy $W^n(\mathbf{x}_1,\mathbf{x}_2) = (\mathbf{y}_{3,1},\mathbf{y}_{3,2})$. Here $(\mathbf{y}_{3,1},\mathbf{y}_{3,2})$ is good iff $\mathbf{y}_{3,1}$ is good. As $\mathbf{y}_{3,2} = \mathbf{x}_2$, we analyze the number of possible codewords $\mathbf{x}_1$ that satisfy $W^n(\mathbf{x}_1,\mathbf{x}_2) = (\mathbf{y}_{3,1},\mathbf{x}_2)$.

Claim 3.2: For any sufficiently large n, with high probability over our code design, for any good $\mathbf{y}_{3,1}$, there are at most $\lceil 8/\delta \rceil$ codewords $\mathbf{x}_1 \in \{\mathbf{x}_1(w)\}_w \cup \{\bar{\mathbf{x}}_1(w)\}_w$ such that $W^n(\mathbf{x}_1,\mathbf{x}_2) = (\mathbf{y}_{3,1},\mathbf{x}_2)$.

Proof: Let $\mathbf{y}_{3,1}$ be good. Recall that the codebook $\{\mathbf{x}_1(w)\}_w$ is chosen uniformly from subsets of size $|\mathcal{M}_1^2|$ in $\mathcal{X}_1^n$. Consider choosing the codebook $\{\mathbf{x}_1(w)\}_w$ in an iterative manner, where in iteration w, $\mathbf{x}_1(w)$ is chosen uniformly from $\mathcal{X}_1^n \setminus \{\mathbf{x}_1(w')\}_{w'<w}$, i.e., the set $\{\mathbf{x}_1(w)\}_w$ is chosen uniformly without repetitions. For any w, the probability (over the choice of $\mathbf{x}_1(w)$) that $W_1^n(\mathbf{x}_1(w),\mathbf{x}_2) = \mathbf{y}_{3,1}$ is exactly the probability that $\mathbf{x}_1(w) \in \mathbf{X}_1(\mathbf{y}_{3,1})$ which is at most $\frac{2^{n/2}}{4^n - 2^{3n/2-\delta n}} \leq 2 \cdot 2^{-3n/2}$ for sufficiently large n. For any w, the probability that $W_1^n(\bar{\mathbf{x}}_1(w),\mathbf{x}_2) = \mathbf{y}_{3,1}$ is exactly the probability that $\bar{\mathbf{x}}_1(w) \in \mathbf{X}_1(\mathbf{y}_{3,1})$ which is again at most $2 \cdot 2^{-3n/2}$. Thus the probability (over the choice of $\mathbf{x}_1(w)$) that $W_1^n(\mathbf{x}_1(w),\mathbf{x}_2) = \mathbf{y}_{3,1}$ or $W_1^n(\bar{\mathbf{x}}_1(w),\mathbf{x}_2) = \mathbf{y}_{3,1}$ is at most $4 \cdot 2^{-3n/2}$. Moreover, the probability that there exist ℓ messages $\{w_1,\ldots,w_\ell\}$ in $\mathcal{M}_1^2$ such that for all $i = 1,\ldots,\ell$ it holds that either $W_1^n(\mathbf{x}_1(w_i),\mathbf{x}_2) = \mathbf{y}_{3,1}$ or $W_X^n(\bar{\mathbf{x}}_1(w_i),\mathbf{x}_2) = \mathbf{y}_{3,1}$ is at most

$$\begin{aligned}\binom{2^{3n/2-\delta n}}{\ell} 4^\ell 2^{-3\ell n/2} &\leq 2^{3n\ell/2-\delta n\ell} 4^\ell 2^{-3\ell n/2} \\ &= 4^\ell 2^{-\delta n \ell}\end{aligned}$$

Setting $\ell = \lceil 8/\delta \rceil + 1$, we have that the above probability is at most 2^{-7n} for sufficiently large n. Union bounding over all 2^{6n} possible good $\mathbf{y}_{3,1}$, we conclude the assertion with probability $1 - 2^{-n}$ over code design. ■

Assume a codebook $\{\mathbf{x}_1(w)\}_w$ that satisfied the conditions of Claim 3.2. We now complete the description of our encoding and decoding scheme of phase 2. Let $\mathbf{x}_1(w)$, $\mathbf{x}_2$ be the codewords forwarded to the CF during the first phase of communication and let b be the one bit value sent back to transmitter node 1 (from the CF) at time step $n+1$. If $b = 0$, let $\mathbf{y}_{3,1} = W_1^n(\mathbf{x}_1(w),\mathbf{x}_2)$ otherwise let $\mathbf{y}_{3,1} = W_1^n(\bar{\mathbf{x}}_1(w),\mathbf{x}_2)$. The vector $\mathbf{y}_{3,1}$ is received at receiver node 3 during phase 2. We are guaranteed that $\mathbf{y}_{3,1}$ is good. By Claim 3.2 we know that at most $\lceil 8/\delta \rceil$ other codewords $\mathbf{x}_1$ may also satisfy $\mathbf{y}_{3,1} = W_1^n(\mathbf{x}_1,\mathbf{x}_2)$. Thus receiver node 3 can decode to a list $L(\mathbf{y}_{3,1})$ of codewords (i.e., messages) corresponding to $\mathbf{y}_{3,1}$ of size at most $\lceil 8/\delta \rceil$ that includes the codeword $\mathbf{x}_1(w)$ or $\bar{\mathbf{x}}_1(w)$ sent by transmitter node 1.

To decode the message M_1^2 correctly, receiver node 3 needs to know whether $\mathbf{x}_1(M_1^2)$ or $\bar{\mathbf{x}}_1(M_1^2)$ was sent by node 1 and in addition the index of $\mathbf{x}_1(M_1^2)$ or $\bar{\mathbf{x}}_1(M_1^2)$ in the list $L(\mathbf{y}_{3,1})$. Here we use the fact that the codebook $\{\mathbf{x}_1(w)\}_w$ consists of distinct codewords. This information can be sent from the CF to receiver node 3 on the k, 1-bit channels of $\mathcal{N}^+$. Here we have that $k = \lceil \log(8/\delta) \rceil + 1$. Specifically, at time step $n+1$, the CF can simulate the behavior of the channel W^n (as W is deterministic) and thus deduce $\mathbf{y}_{3,1}$ and the list $L(\mathbf{y}_{3,1})$ available to receiver node 3. Using a lexicographic ordering on all vectors of $\mathcal{Y}_{3,1}^n$, at time step $n+1$, the CF can send the bit b (the same bit it sends to transmitter node 1) and also the location of $\mathbf{x}_1(M_1^2)$ or $\bar{\mathbf{x}}_1(M_1^2)$ in $L(\mathbf{y}_{3,1})$ to receiver node 3. This information allows node 3 to determine M_1^2 (without error). This concludes the proof for phase 2 of our communication.

B. Putting it all together

All in all, in $2n+1$ time steps, the message $M_1 = (M_1^1, M_1^2) \in [2^n] \times [2^{\frac{3n}{2}-\delta n}]$, and the message $M_2 = (M_2^1, M_2^2) \in [2^n] \times [2^n]$ are decoded with zero error at receiver node 3. This implies a communication rate of

$$\left(\frac{2n(1.25-0.5\delta)}{2n+1}, \frac{2n}{2n+1}\right)$$

which in turn implies for $\delta = 0.02$ that

$$(1.24, 1) \in \mathcal{C}_{\max}(\mathcal{N}^+)$$

as asserted in the theorem.

IV. Conclusions

In this work we study the effect of 1-bit of communication in memoryless networks. We prove the existence of a network $\mathcal{N}$ and a 1-bit channel e such that adding e to $\mathcal{N}$ yields a new network $\mathcal{N}^*$ with a strictly larger capacity region. Our results hold for the maximal-error criteria. Whether there exist a memoryless network $\mathcal{N}$ and a δ capacity noiseless channel e (let alone a 1-bit channel e) such that adding e to $\mathcal{N}$ results in a capacity region $\mathcal{C}_{avg}(\mathcal{N}(\delta))$ which is not continuous at $\delta = 0$ is an intriguing open problem. The effect of adding a 1-bit channel in the noiseless setting of network coding is not addressed in this work and remains a fascinating open problem.

References

[1] T. Ho, M. Effros, and S. Jalali. On equivalences between network topologies. In *Forty-Eighth Annual Allerton Conference on Communication, Control, and Computing*, 2010.

[2] S. Jalali, M. Effros, and T. Ho. On the impact of a single edge on the network coding capacity. In *Information Theory and Applications Workshop (ITA)*, 2011.

[3] V. Guruswami. List decoding with side information. In *Proceedings of 18th IEEE Annual Conference on Computational Complexity*, pages 300–309, 2003.

[4] M. Langberg. Private codes or succinct random codes that are (almost) perfect. In *Annual IEEE Symposium on Foundations of Computer Science*, volume 45, pages 325–334, 2004.

[5] P. Noorzad, M. Effros, and M. Langberg. Can negligible cooperation increase network reliability? In *International Symposium on Information Theory*, 2016.

[6] G. Dueck. Maximal error capacity regions are smaller than average error capacity regions for multi-user channels. *Probl. Contr. Inform. Theory*, 7:11–19, 1978.

[7] T. Chan and A. Grant. On capacity regions of non-multicast networks. In *International Symposium on Information Theory*, pages 2378 – 2382, 2010.

[8] T. Chan and A. Grant. Network coding capacity regions via entropy functions. *IEEE Transactions on Information Theory*, 60(9):5347–5374, 2014.

[9] M. Langberg and M. Effros. Network coding: Is zero error always possible? In *Proceedings of Forty-Ninth Annual Allerton Conference on Communication, Control, and Computing*, pages 1478–1485, 2011.

[10] M. Langberg and M. Effros. Source coding for dependent sources. *In proceedings of IEEE Information Theory Workshop (ITW)*, 2012.

[11] M. F. Wong, M. Langberg, and M. Effros. On a Capacity Equivalence between Network and Index Coding and the Edge Removal Problem. In *proceedings of International Symposium on Information Theory*, pages 972 – 976, 2013.

[12] M. F. Wong, M. Langberg, and M. Effros. On an equivalence of the reduction of k-unicast to 2-unicast capacity and the edge removal property. In *IEEE International Symposium on Information Theory (ISIT)*, pages 371–375, 2015.

[13] M. F. Wong, M. Langberg, and M. Effros. On the tightness of an entropic region outer bound for network coding and the edge removal property. *To appear in IEEE International Symposium on Information Theory*, 2016.